AD-A183 942 PRACTICAL APPLICATIONS OF ARTIFICIAL INTELLIGENCE EXPERT SYSTEMS AT NAVSEA(U) ASSOCIATION OF SCIENTISTS AND ENGINEERS OF THE NAVAL SEA SYST.. P J HARTMAN 1/1

UNCLASSIFIED 03 APR 87 F/G 12/9 NL

END
9-87
DTIC

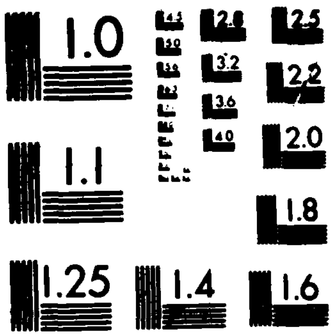

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963-A

FILE COPY

24TH ANNUAL TECHNICAL SYMPOSIUM 1987

"NAVY IN THE NINETIES — TREND AND TECHNOLOGY"

AD-A183 942

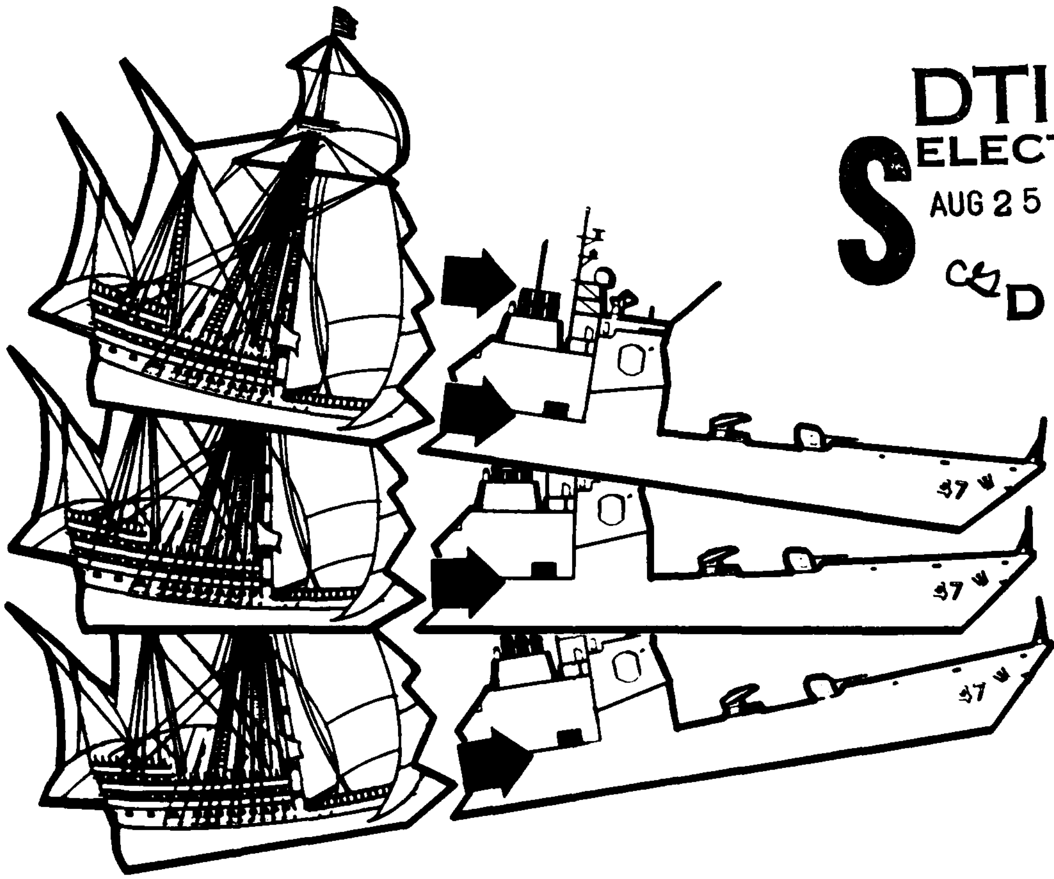

PRACTICAL APPLICATIONS OF ARTIFICIAL INTELLIGENCE, EXPERT SYSTEMS AT NAVSEA

by: Patrick J. Hartman

ASSOCIATION OF SCIENTISTS AND ENGINEERS OF THE NAVAL SEA SYSTEMS COMMAND • DEPARTMENT OF THE NAVY • WASHINGTON, D.C. 20362

Practical Applications of Artificial Intelligence

Expert Systems at NAVSEA

Patrick J. Hartman, Ph.D., P.E.

Head, Design Engineering Section

Reliability and Maintainability Engineering Branch

Materials and Assurance Engineering Office

3 April 1987

Approved for Public Release
Distribution Unlimited

The views expressed herein are the personal opinions of the author and are not necessarily the official views of the Department of Defense or of the Department of the Navy

ABSTRACT

This paper examines the historical oscillations of Artificial Intelligence (AI). It traces the very beginnings from 1956 at the Dartmouth Summer Research Project on Artificial Intelligence, through the depressions of the 1960's and 1970's, to the exponentially increasing number of successful applications in the mid 1980's. It shows how to select problems which can be solved with the aid of Expert Systems; it ~~The paper~~ also identifies areas such as common sense reasoning, automated machine learning, and complex design synthesis which are beyond the state of the art now and will be for years to come.

Artificial Intelligence programs are able to diagnose faults and classify solutions in narrowly defined specialties even when the data is "fuzzy", but they have not exhibited autonomous "thinking". Just as conventional computer programming has alleviated the burden of calculating, AI Expert Systems will streamline the processing of logical data. Both of these computer techniques are cost effective when they are applied to well defined tasks since computers are faster than people and error free for routine tasks.

Expert Systems are interesting and useful tools, but they are not magic. They require significant development effort, but with them we can perform tasks which were not possible before. These programs are fundamentally different than those which performed numerical calculations in the past. They process information - the rules by which equipment and people act. Certain situations will always occur which require a "real expert", but every engineer can use an assistant and they are being developed.

Table of Contents

List of Figures

List of Tables

I. What is Artificial Intelligence and What Can it Do Now

What is Artificial Intelligence (AI)? Does anyone really know? The literature is full of conflicting definitions. Dreyfus (January, 1986) indicates that 25 years ago researchers dreamed of "...autonomous "thinking" machines that are free of human control. And now they believe we are not far from realizing that dream." Well, we may be close but I am not so sure and neither is Dreyfus.

Other concepts include Daniels' (June, 1986) "AI is defined as the application of knowledge, thought and learning to computer systems to aid humans." Whatever the definition is, the important aspects of artificial intelligence are that this is a new field of computer implementation which manipulates knowledge and symbols in ways which were not possible with conventional data processing. (Robinson, March 1986)

In AI, the latest information appears to be the best. The same appears to be true for the commercial software being developed for AI. Each new product builds on the successes of the previous. This is why you will find over 80 references at the end of this paper. These are taken from the current literature from dozens of readily available sources and can be used to guide anyone in their examination of AI.

Subfields Within Artificial Intelligence

Artificial Intelligence could be the next great innovation in science and engineering, but the field is very broad (Table 1). AI includes applications in:

- Artificial Vision (Reitman, 1983),
- Expert Systems,
- Information Fusion,
- Machine Learning,
- Management (Gautschi, 1986),
- Natural Language Interfaces (Strickler, September 1986),
- Programming Tools,
- Robotics (Marling, March 1986)
- Voice and Sound (USA TODAY, Wood 1986).

Not every one of these areas is applicable to NAVSEA. Not all of them are practical with today's AI technology. This paper will focus only on the practical tools which can be applied to our design and management functions. We will concentrate on Expert Systems (Knowledge Based Systems).

What are Expert Systems?

Expert systems are methods and programs used to solve complex problems involving information, rules of thumb, and uncertainty. They can be used to transport knowledge to non-experts and also assist real experts by performing rote tasks expertly. These expert systems can be very useful to us as long as we concentrate on what is practical today. (Head, August 1986, October 1986).

These systems can be used to store knowledge as it is developed, preserve the knowledge of people retiring from the organization, provide advice given uncertain inputs, quickly search mountains of data, and pick up details often skipped over by people in a hurry. (See Fersko-Weiss, November 1985, Menke, August 1986, Sullivan, August 1986).

But the expert system based on artificial intelligence does not reason the same way a person does. It does not perceive significance nor jump intuitively to the conclusion. They are narrow in their understanding. Right now the expert systems are more like apprentices, but they are capable of solving complex problems.

According to Hayes-Roth (October 1984), The limit of human expertise is about 100,000 rules. Expertise in a profession requires about 10,000 rules. This can be cut down to between 500 and 1000 for expert competence in a narrow area. (See Table 2.) This narrow expert, which is often refered to as a knowledge based system, can be of real use to our engineers, naval architects, and managers.

Hayes-Roth (1984) states, "An expert system is a knowledge-intensive program that solves problems normally requiring human expertise. It performs many of the secondary functions that an expert does, such as asking relevant questions and explaining its reasoning." He also says that they do (or are supposed to do) the following:

- "They solve very difficult problems as well as or better than human experts,
- They reason heuristically, using what experts consider effective rules of thumb,

- They interact with humans in appropriate ways, including the use of natural language,
- They manipulate and reason about symbolic descriptions,
- They function with erroneous data and uncertain judgmental rules,
- They contemplate multiple hypotheses simultaneously,
- They explain why they're asking a question, and
- They justify their conclusions."

The Boom/Bust Cycles of AI

John McCarthy, inventor of LISP, coined the term 'Artificial Intelligence' for the 1956 summer workshop at Dartmouth College (D'Ambrosio, 1985, Kumara, 1986, Waldrop, 1984). Since then, AI has had at least two periods of heightened expectations and crushed hopes. The late 1950s and early 1960s focused on autonomous thinking, but it hasn't come yet (Dreyfus, 1986). D'Ambrosio (1985) finds that "During the 1960s, computer scientists developed a number of general problem-solving mechanisms, and in the late 1960s and early 1970s they tried to apply these mechanisms to "real" problems. For the most part these attempts resulted in dismal failure. In fact the results were so disappointing that one country, Great Britain, completely abandoned its AI research and development effort."

Some people said AI could do anything. That is absolutely untrue. It is just another computer tool. It is not magic. (Schoen, 1985) Twenty years ago, when the first medical and mathematical expert systems were being written, it took about forty staff-years of effort to get them working according to Winston (1984).

This situation has changed dramatically over the past five years. Generic knowledge systems which embody natural language interfaces, tools for developing the expert knowledge base through rules and examples, and the inference engines are now commercially available for desk top micro-computers.

It appears that AI has achieved 'critical mass' and that the field will continue to grow and provide practical applications.

"IBM plans to use AI internally to strengthen its competitive position." (Heffernan, September 1986). The Digital Equipment Corporation (DEC) "...has about 400 engineers working full time on artificial intelligence products or internal tools in Hudson, Mass., Tokyo and Valbonne, France". (Brodie, October 1986). Other major computer companies are spending about $30-40 million per year on AI. With this kind of commercial effort AI looks solid. (Rauch-Hindi, November 1985, August 1986, Engineering Times, October 1986)

Major Money Making/Saving Areas

As stated before, AI can not do everything (Milne, October 1985), but there are areas like Problem Classification, Diagnosis and Correction, Situation Analysis and Understanding, and Design Synthesis where the tools are applicable and cost effective. Chandrasekaran (November 1986, Reitman, 1984) has taken considerable effort to determine which AI tools can successfully solve different problems and which can't. Table 3 synopsizes his work and expands on it. Classification problems, the Structured Selection from a hierarchical structure of possible solutions is the most successful. This technique forms the basis of today's expert systems.

When an airline sells one additional seat on each flight by using an expert system to assist in routing and pricing, they can make an additional $12 million per year. The Defense Advanced Research Project Agency (DARPA) is spending a lot of money in hopes that military expert systems can aid battlefield tactics (Design News, October 1985). DARPA is presently spending $150 million per year (Fields, 1985) on expert system and natural language interface development and the Japanese have targeted their Fifth Generation computers for these areas. The FBI has implemented Little Floyd a "...fledgling expert system...that draws on established data bases to produce...inferences...on suspicious business transactions, unusual contacts and other telltale signs in relationships among crime figures..." (Head, October 1986, Schrage, July 1986).

The Digital Equipment Corporation says it saves $18 million per year using its expert system to configure new VAX computer sites (Rausch-Hindi, October 1986).

The Internal Revenue Service is reviewing proposals submitted to "develop a knowledge-based expert-system prototype to automate the current manual process of identifying issues with a high probability of adjustment on individual 1040-type tax returns." (Sullivan, July 1986).

This area of classification also extends to Diagnosis and Correction, maintenance trouble shooting on jet aircraft (Design News, July 1986) and Intelligent Service Manuals in which week long problem diagnoses can be performed in five minutes using an expert system. Expert advice on ship routing based on changing weather conditions also appears feasible (MARAD, June 1986).

Ship routing may be an example of Situation Analysis and Understanding in much the same way that acoustic signature analysis (Groundwater, 1985), crisis management, information fusion, and oil dirlling analyses (Sea Technology, June 1986) are. These are all structured classification.

II. How to Pick and Implement the 'Right' Problems

It takes considerable time for an organization to develop its first successful expert system. The 'right' type of problems have to be selected and a corporate infrastructure must be established to initiate the development. NAVSEA is doing this through the Computer Supported Design Division's Working Group for Artificial Intelligence.

These problems can be divided into two basic catagories.

- Corporate Level Problems which nothing less than AI can resolve. These can cost millions to develop and implement.

- Personal Level day-to-day problems which Expert Systems can assist in solving and save many staff-years when used. These cost between $35,000 and $70,000 to develop and implement.

Either way the savings can be great. The corporate level developments are intense and the cost is significant. Personal Level Knowledge Based Expert Systems are distributed over a wide range of needs.

There are several problem catagories, as stated before. Classification problems have been described as well developed and highly successful. Table 3 also shows Design Synthesis to be successful now, but correlating abstract data from a complex data base is still under development. This capability could be available soon (Moad, 1987).

The remaining AI applications on Table 3 are not practical yet and should be avoided for engineers and naval architects interested in cost effective results. This is particularly true of systems which need to rely on 'common sense'. Our 'common sense' is based on years of experience and millions of facts and tens of thousands of rules. The present software and hardware are not capable of this level and probably won't be for another AI generation (possibly a decade).

Scoping the Problem

Numerous expert systems have been implemented on large computers at great cost during the 1960s and 1970s by researchers in the medical and geological fields. Their work has led to a compact list of criteria for developing commercially feasible expert systems. Hayes-Roth and Winston (1984) have explored the range. The rules applicable to the NAVSEA effort are synopsized and expanded below:

- Select high value problems,

- Choose an area where there are recognized experts. Experts that have a better track record than the amateurs. Winston says that this rules out astrology and probably selecting stock portfolios,

- Rule out those problems where the experts always argue about whether the answer is right or not,

- Pick problems that an expert could solve over the phone. This insures that the problem is computer compatible and eases the data interface,

- Find problems that typically take an expert between a few minutes and a week to solve. Anything shorter leaves little to be gained by computerizing and anything longer is most likely beyond the state of the art for the next few years,

- Pick a problem in which the skills are routinely taught to "neophytes",

- Select the problem and ask the "expert" to review it,

- Develop a knowledge base that contains the rules used by the expert to solve the problem. This information feeds the inference engine which makes decisions with respect to the data provided to the expert system,

- Have the computer program solve several training problems,

- Have the real expert critique the computer "expert" and make recommendations,

- Improve the expert system by incorportating the expert's critique, and

- Continue to train the expert system and augment its knowledge base incrementally.

Rule Development

Twenty years ago it took twenty to forty staff-years to develop expert systems. A Design Synthesis expert system used to configure computer systems took ten staff-years to develop in 1978. Winston (1984) has done a fine job in defining this productivity trend and Hayes-Roth (1984) plots a "doubling of productivity every two years over the last 15 years."

Hayes-Roth's graph showed that 500 rule expert systems could be produced with about one staff year of effort using commercially available tools sold in 1982-83 (Schlosberg, May 1985). Hayes-Roth anticipated that generic knowledge systems would be sold in 1984 or 1985. Tucker showed that they were commercially available for microcomputers in 1985. With genaric knowledge systems, coding 500 rule expert systems can take as little as 500 hours.

Hayes-Roth's curve has been reproduced in Figure 1 and the information expanded upon. Programming languages are becomming more efficient and Dourson (November 1986) indicates that by 1990 people will be able to implement one rule per day with this technique, but this is still almost an order of magnitude slower than implementation with today's expert system shells. With today's commercially available expert systems shells, it does take about one hour to develop and implement each rule. This number has been confirmed by Vesonder (November 1985) who developed a 50 rule system with his experts over one weekend. This is phenomenally better than the 2000 staff hours needed to develop and implement each rule in a programming language like LISP during the late 1960s. It was like reinventing expert systems for each new implementation.

One hour is approximately the time needed for a person to develop a useful rule. This means we are no longer tool speed limited in our development but staff limited. In order to break this barrier, automated machine learning is needed (Hanson, November 1986). Machine learning is still a research and development area (Kolokouris, November 1986, Thompson, November 1986).

Expert System Software Characteristics

As part of its effort to examine the practical application of AI and expert systems for NAVSEA, the AI Working Group has investigated over forty of the currently available commercial expert system programs. The capabilities of these tools vary greatly as do their prices. The AI Working Group purchased and tested five of these packages. The price of the software purchased ranged from $95.00 to $1800. We did not purchase the highest priced tool which cost $85,000.

Price does not necessarily correlate with the capability of the tool. In general the higher the price, the greater the number of features included, but currency is also very important because each new tool builds on the latest advances in the field.

Truly useful expert systems include the following characteristics:

- Plain English Programming - the Natural Language Interface so necessary from the human factors standpoint,
- Rules and Frames, the sets of rules which can be attached to entities including objects. These rules form the Knowledge Base from which the results are derived,
- An Inference engine, with both forward and backward chaining (searching), which draws conclusions from the Knowledge Base,
- Sometimes a Knowledge Engineer is needed to help the real expert derive rules for the knowledge based system,
- The ability to handle Uncertainty in the input and in the knowledge base,
- The ability to place Confidence Factors on the results from this uncertain information,
- Give Explainations as to how the program arrived at the advice given,
- Provide mathematical and statistical Functions within the program to facilitate and enhance the rules,
- Provide standardized computer Interfaces to other programs and large data bases so that sophisticated tasks can be accomplished,

- Be written in a standard, transportable Language such as LISP or C so that it can be used on many machines without rewritting and for operational speed on the development and general purpose delivery machines chosen, and

- Have a Price that is reasonable for the task being undertaken.

Of the expert systems surveyed, all of them used rules for the knowledge base. Some, As shown in Figure 2, were able to process frames of information also. This capability was included only in those programs costing over $5000.

Software priced from $95 to $5000 typically runs on small personal computers. From $5000 to $85,000 the development programs fit better on higher powered mini and mainframe computers. Some of the most capable programs can operate efficiently over the entire range from micro to mainframe environments without the need for specialized machines. These programs are typically written in C for speed of operation on general purpose machines which do not have fast LISP capabilities. (Verity, August 1986)

Figure 2 compares the characteristics of these two price classes of expert system shells. Note that specially trained knowledge engineers are typically needed to help write the applications for the highest priced software because of their complexity and widely ranging capabilities.

Table 4 defines the initial software capabilities needed to implement practical ship design and managerial knowledge based systems at NAVSEA. We will rapidly outstrip the minimum 1000 rule set. Frame capability and 10,000 rules will probably be needed in the next year or so. Much depends on the size and complexity of the problems chosen.

Floors and Ceilings

In general, the more expensive expert systems have greater capabilities, but they are also harder to use because of their current complexity. Figure 3 qualitatively illustrates this phenomenon. The small personal computer based expert systems may be easier to learn, have a low floor, but thier ultimate capability (ceiling) is lower than the more expensive tools. The true challenge is to determine which expert system from this spectrum of capability and ease of use matches the problem to be solved.

Interfaces and Input/Output

In the past (Vesonder, November 1985), about 20% of the total expert system development effort using the primary languages in artificial intelligence went into deriving and implementing the rules. The rest of the time was used for interfacing with other programs, the user, and providing understandable output. It is hoped that the new expert system shells can turn this around to be 80/20 in favor of the knowledge based efforts.

The newest software tries to alleviate this interfacing problem by integrating several user oriented input and output capabilities in the same package as the expert system (Heilliwell, May 1986, Holsapple, 1986). The software includes:

- The Expert System,
- A Natural Language Interface,
- A Relational Data Base,
- Numerical Functions,
- Statistical Analysis,
- Spreadsheets,
- Business Graphics.
- Text Processing,
- Forms and Reports,
- Remote Communications and Interfacing with Other Programs and Data Bases.

In the final judgement, program capability and ease of use are very important, but the knowledge placed in the expert system is most critical.

Hardware Needs

Which development and delivery computers are needed to support engineering design and management expert systems? Rauch-Hindin, (November 1985) indicates that the workstation environment is well suited to this level of task. This reinforces our benchmark findings in Figures 4 and 5, that any processing speed is possible, but that current NAVSEA computer hardware is capable of running AI problems as fast as dedicated LISP machines. Vector supercomputers are faster, personal computers are slower, and the mini-computer, workstations are about right for processing AI and conventional programs.

The balance between general purpose mini-computers and special LISP machines is shown in Figure 5. The baseline NAVSEA mini-computer is able to perform satisfactorily in the AI LISP environment. Specialized LISP machines may not be able to perform FORTRAN calculations as fast as general purpose mini-computers because their compilers have not necessarily been optimized for that language (Verity, August 1986).

III. The Costs and Effort Required to Develop an Expert System

There are large and small problems to be solved cost effectively with expert systems. The large ones are called "Gold Nuggets" (Feigenbaum, June 1986) because of the potential and demonstrated payoffs achieved by successful expert systems in these areas. Failure prediction of main power turbines in the electrical generation industry is one example of multi-million dollars savings from expert systems.

Large, corporate level problems also take hundreds of thousands or millions of dollars to implement. Table 5 illustrates the typical cost breakdowns for developing and using corporate and personal level knowledge based systems. This data was developed from the cost and staffing data from numerous successful expert systems derived from current AI languages (Buchanan 1985, Davis 1985, Vesonder 1985, Feigenbaum 1986, Mahler 1986). Development using Expert Systems Shells should be somewhat less costly.

Each system required the hardware, development, and maintenance efforts shown below:

- Hardware. Large, specialized, dedicated LISP machines for the corporate level. Personal computers and workstations for the rest,
- Software,
- Problem Assessment. About one staff-month to pick the right problems and lay out the overall program,
- Feasibility Study. Two Staff months for large systems,
- Prototype System Development. One to two staff-years for corporate level expert systems. One to three staff-months for small systems,
- Product Development. Development of the production knowledge based system for large problems, one to ten staff years in one to four calendar years,
- Benchmarking. Three staff-months for large problems,.
- Training, and
- Yearly Maintenance. Very significant for corporate level programs. Performed by the user of personal level expert systems.

While the costs of the large and small systems were vastly different, the proportion of costs going into each of the above areas were very similar. Table 5 details the percentage of costs in each area and Figure 6 graphically demonstrates the similarities between the two sizes of systems.

Hardware costs are 10 percent and software 5-10 percent. The major costs are prototype and production development. This indicates that spending the upfront money for hardware, software, feasibility studies, and problem assessment are a very cost effective way of avoiding expensive problems in the later phases of expert system development.

IV. Areas of Potential NAVSEA Application

Expert systems can be applied at NAVSEA on a cost effective basis right now. Several general areas of interest are listed below. These were areas identified by the AI Working Group and are not intended to be an all inclusive list. When a specific problem needs to be solved, always consult the area classification guide shown in Table 3.

- Budget Planr ng, (Barnes, November 1986)
- Ballast Control, (Williams, July 1986)
- Computer Aided Design and Computer Aided Engineering (CAD/CAE),
- Damage Control,
- Equipment Status Monitoring on Surface Ships and Submarines,
- Fire Fighting,
- Intelligent Service Manuals,
- Management and Program Control,
- Manning Analysis,
- Paint Failure Analysis,
- Reliability Analysis,
- Shipboard Arrangements,
- Shipboard Equipment Configuration Control,
- Ship Alts,
- Ship Specifications,
- Shock Qualification Extensions,
- Technical Training, and
- Welding Problem Analysis.

Several of these areas are being explored and developed by the technical societies and institutes. The information gained by them can and should be used in the NAVSEA development effort. Through the careful use of tailored commercial software and knowledge gained by other expert sources, we can rapidly bring this new technology to bear on our problems.

V. What Developments Can Be Expected Next

If expert systems can be cost effectively implemented now, what can we expect in the future? The field of Artificial Intelligence is a rapidly changing one. AI could go bust again, but that seems unlikely. The probable path is that various aspects of AI will be included in more and more software during the next one to five years. It may become a transparent part of more 'user friendly' software.

Again, the latest is probably the greatest and we may be seeing integrated CAD/CAE, expert systems, data bases, and hardware all operating together in the workstation environment within the next five years.

Several advances are in development right now:

- Increased Speed on Conventional Machines. Better LISP and C optimization is increasing the development and run speed of expert systems on all machines. (Robinson, January 1986, Serlin, 1986, Verity, August 1986),

- Supercomputers on a Chip. Very High Speed Integrated Circuit (VHSIC) technology makes this development possible. These ultrapowerful chips containing about 35 million transistors will make almost instantaneous processing of expert systems possible on a workstation. (Newport, September 1986, Norton, October 1986),

- Parallel Processing. AI is one of the driving forces behind the development of new computers which have over 64,000 processors acting in parallel. By working on various parts of a problem at the same time, these machines may be able to reach solutions faster. The immediate challenge is to cast the real world problems in terms which can be processed by these computers. (Datamation, July 1986), and

- Neural Networks. Research is proceeding in the effort to use Metal Oxide Simiconductor Field-Effect Transistors (MOSFET) technology to mirror the way simple biological neural networks act. (Hopfield, August 1986, Tucker, July 1986, Victor, August 1986) If these experiments work, it may be possible to develop 'thinking machines', but it may not be for several years.

The only advice that can be given in this situation is; work with the best technology available and plan for the future enhancements.

VI. Conclusions

Expert systems are practical. Programs are commercially available which can be tailored to match our needs and solve our problems. The cautions are that AI can't do everything and that we need to be very careful in applying expert systems to insure success of the projects. Knowledge acquisition takes a significant amount of work and we need to intentionally dedicate the time of our real experts to the task.

NAVSEA needs the best software to solve practical problems. The expert system software must support natural language interfaces, rules and frames as input to the knowledge base, data base interfaces, the ability to execute other programs and retrieve the information, and give explainations of how the system arrived at the answer.

It is up to NAVSEA engineers, naval architects, and managers to:

- Examine the field of expert systems,
- Get people involved,
- Pick the most promising NAVSEA problems, document them, and distill the expert knowledge required to solve them,
- Make this technology available,
- Build, test, and use the systems to save the Navy time and money, now and in the future.

The development will not be easy, training is required to understand what expert systems can do and how to go about building the programs from tailored commercial shells.

It is time to think creatively and build cautiously. We need to pick feasible problems, build a cadre of knowledgable people in numerous NAVSEA groups, and produce cost effective expert systems to assist our engineers and managers.

VII. References and Bibliography

Allen, L., The Cost of an Expert, Computerworld, July 21, 1986 pp. 64-68

Barnes, R., Hosack, G., Tonn, B., and Kunz, D., Expert Systems and the Navy Budgeting Process, Oak Ridge National Laboratory, Oak Ridge, TN, November 6, 1986, pp. 1-12

Brodie, R., "For 5th Generation, the Language Just May Be Japanese," Digital News, October 20, 1986, p. 22

Brown, M., "Expert Systems Pose Liability Problem in Courts," Government Computer News, August 15, 1986, p. 20

Buchanan, B., The First Artificial Intelligence Satellite Symposium, Texas Instruments, November 13, 1985

Chandrasekaran, B., "Generic Tasks in Knowledge-Based Reasoning," U.S. Navy Center for Applied Research in Artificial Intelligence Seminar, November 3, 1986, pp. 1-4

Conlon, T., "Expert Systems in the Economy Size", Computer Decisions, May 21, 1985, pp.32-33

D'Ambrosio, B., "Expert Systems-Myth or Reality?'" BYTE, January 1985, pp. 275-282

Daniels, J. D., "Artificial Intelligence: A Brief Tutorial," Signal, June, 1986, pp. 23-26

Datamation, "Multimicro Super," July 1, 1986, p. 96

Davis, R., The First Artificial Intelligence Satellite Symposium, Texas Instruments, November 13,1985

Design News, "Military Expert System to Aid Battlefield Tactics," October 21, 1985, pp. 21-22

Design News, "Jet Fighter uses AI as Troubleshooter," July 7, 1986, p. 20

Design News, "Computer Incorporates 64,000 Processors," July 21, 1986, p. 19

DiDio, L., "VAX 8700 Performance Benchmarks Show DEC's Claims Were Right: It's a Fast Machine," Digital Review, November 10, 1986, p. 8

Digital Review, "Inside AI," May 1986, p. 17

Digital Review, "He's Seen The Light, Artificially Speaking," November 17, 1986, p. 17

Dourson, S., "Knowledge Systems," DEC Professional, November 1986, pp. 46-50

Dreyfus, H., and Dreyfus, S., "Why Computers May Never Think Like People," Technology Review, January 1986, pp. 44-61

Engineering Times, "IBM Interest Gives Artificial Intelligence New Respectability," October 1986, p. 8

Feigenbaum, E., "Knowledge - Based Systems: A Step-By-Step Guide to Getting Started," The Second Artificial Intelligence Satellite Symposium, Proceedings/Source Book, Texas Instruments, June 25, 1986.

Fersko-Weiss, H., "Expert Systems Decision-Making Power," Personal Computing, November 1985, pp. 97-105

Fields, C., "DARPA Strategic Computing Program," 24th Annual Technical Symposium of the Washington, D.C., Chapter of the ACM, June 20 1985

Fredell, E., "AI Agent Nabs Crooks," Government Computer News, June 6, 1986, p.19

Gautschi, Dr. T.F., P.E., "A Management Perspective on Artificial Intelligence," Design News, May 5, 1986, pp. 174-175

Groundwater, E. H., "An Expert System for Ocean Surveillance," IEEE, January 1985, pp. 196-199

Hayes-Roth, F.,"Knowledge Based Expert Systems," October 1984 pp. 263-273

Hays, N., "AI Programs Permit Development of Expert Systems on PCs, Computer, July 1986,

Head, R. V., "Artificial Intelligence Aficionados Offer Nothing New," Government Computer News, August 15, 1986, pp. 25-26

Head, R.V., "Concentrating on AI Tools That Are Practical Today," Government Computer News, October 10, 1986, p. 33

Hefferman, H., "Exec Reveals IBM AI Directions," Government Computer News, September 12, 1986,p. 35

Helliwell, J., "GURU: Brave New Expert System?," PC Magazine, May 27, 1986, pp. 151-163

Holsapple, C., Whinston, A., "Manager's Guide to Exepert Systems Using Guru," Dow Jones-Irwin, 1986.

Hopfield, J.J., Tank, D.W., "Computing with Neural Circuits: A Model," Science, August 8, 1986, pp. 625-633

Hornstein, J.V., "Parallel Processing Attacks Real-Time World," Mini-Micro Systems, December 1986, pp. 65-77

Jones, J.E., "Developing An Expert System for Engineering," Computers in Mechanical Engineering, November 1986, pp. 10-16

Jupille, T., "Expert System are New Textbooks," Research & Development, December 1986, pp. 52-55

Klinger, D. E., "Rapid Prototyping Revisited," Datamation, October 15, 1986

Kneale, D., "How Coopers & Lybrand Put Expertise Into Its Computers," The Wall Street Journal, November 14, 1986, p.33

Kolokouris, A. T., "Machine Learning," BYTE, November 1986 pp. 225-354

Kumara, S., Soyster, A.L., and Kashyap R.L., "An Introduction to Artificial Intelligence," IE, pp. 9-20

Mahler, E., "Knowledge - Based Systems: A Step-By-Step Guide to getting Started," The Second Artificial Intelligence Satellite Symposium, Proceedings/Source Book, Texas Instruments, June 25, 1986

Margolis, N., "Software Patent Cheered by AI Community," Digital Review, October 13, 1986

Marine Engineering, "Meet the Intelligent Ship," March 1986 p. 53-56

Maritime Administration, "Expert Sailing Directions (A Tool for Climatological Ship Routing)," June 1986, pp. 1-40

Marling, W., "Hands That See," The DEC Professional, March 1986, pp. 29-36

Menke, S. M., "Senate Archivist Sets Up System to Handle Documents," Government Computer News, August 1, 1986 p. 28

Milne, Captain R. W., PhD., "A Few Problems with Expert Systems," pp. 66-67

Moad, J., "Building a Bridge to Expert Systems," Datamation, January 1, 1987, pp.17-19

Moore, A. H., "Computers, Chips, and Factory Automation," Fortune, October 13, 1986, pp. 28-40

Mosley, J.D., "Expert Systems as Engineering Tools will Broaden Productivity/Creativity Options," November 13, 1986, pp. 91-100

Myers, E., "Making Micros Experts," Datamation, September 1, 1985 pp. 50-54

Naecker, P. A., "Intelligent Database Machines," DEC Professional, November 1986

Newport, J. P. Jr., "A Supercomputer on a Single Chip," Fortune, September 29, 1986, pp. 128-129

Pottmyer, J.J., "Intelligent Systems: Their Development and Application," McGregor & Werner, Inc., June 20, 1985

Rausch-Hindin, W., "(AI) Low-Cost Run-Time Vehicles Dominate Commercial AI," Systems & Software, November 1985, p. 24

Rausch-Hindin, W., "AI Shapes Up for Mainstream Use," Mini-Micro Systems, August 1986, pp. 57-68

Rausch-Hindin, W., "Software Integrates AI, Standard Systems," Mini-Micro Systems, October 1986, pp. 69-85

Reitman, W., Artificial Intelligence Applications For Business, Ablex Publishing Corporation, May 1983, p. 63

Robinson, G. M., "Mainframe Technology in a Micro," Design News, January 20, 1986, pp. 119-124

Robinson, G. M., " Artificial Intelligence: An Emerging Technology Whose Time Has Come," Design News, March 3, 1986 pp. 69-156

Rolandi, W.G., "Knowledge Engineering in Practice," AI Expert, December 1986, pp. 58-62

Schlosberg, J., "Almost Human," Digital Review, May 1985, pp. 37-42

Schoen, S., "Artificial Intelligence Minus the Hype: Simply a New Tool," Design News, December 16, 1985, p. 17

Schrage, M., "Big Floyd" All Wired Up To Aid G-Men (U.S. Agencies Using Artificial Intelligence To help Identify Potential Lawbreakers," The Washington Post, July 20, 1986 pp. 1-4

Sea Technology, "Artificial Intelligence Provides Aid for Drilling," June 1986, p. 31

Serlin, O., "MIPS, DHRYSTONES, and other Tales," Datamation, pp. 112-118

Strickler, D., "Exercising the Intellect, in Words You Can Understand," Digital News, September 29, 1986, p. 51

Sullivan, J. A., "IRS Solicits Expert Systems For Tax Effort," Government Computer News, July 18, 1986, p. 15

The Observer, "Computers for Handicapped Employees," October 10,1986, pp. 6-8

Tucker, M. "IBM PC Expert Systems Breed New VAR Markets," Mini-Micro Systems, June 1985 pp. 26-27

Tucker, M., "Coming Next from Japan: The Bionic Computer?," Mini-Micro Systems, July 1986, pp. 28-30

Tucker, M., "Artificial Intelligence Rides Out the Slump," Mini-Micro Systems, September 1986, pp.50-53

USA Today, "20/20 Hindsight," September 22, 1986, p. E1

Verity, J. W., "The Lisp Race Heats Up," Datamation, August 1, 1986, pp. 55-58

Vesonder, G. T., "ACE: An Expert System for Telephone Cable Maintenance," U.S. Navy Center for Applied Research in Artificial Intelligence Seminar, November 18, 1985

Victor, J., "Bell Labs Models Parallel Processor on Neural Networks," Mini-Micro Systems, August 1986, pp. 43-51

Waldrop, M. M., "Artificial Intelligence (I): Into the World," Science, February 24, 1984, pp. 802-805

Waldrop, M. M., "The Necessity of Knowledge," Science, March 1984 pp. 1279-1282

Waterman, D. A., A Guide to Expert Systems, Addison-Wesley Publishing Company, October 1985, pp. xi-339

Williams, D., "Artificial Intelligence Helps Right Ship Stability Control Operations," Defense News, July 21, 1986, p. 13

Williams, R., "Hammer Reminisces About Early Days of UNIVAC," Government Computer News, July 18, 1986, pp. 8-10

Winston, P.H. and Pendergast, K.A., The AI Business - The Commercial Uses of Artificial Intelligence, The MIT Press, Cambridge, MA 1984

Wood, L., "Voice Recognition Research Yields Host of Products," Government Computer News, May 23, 1986, p. 80

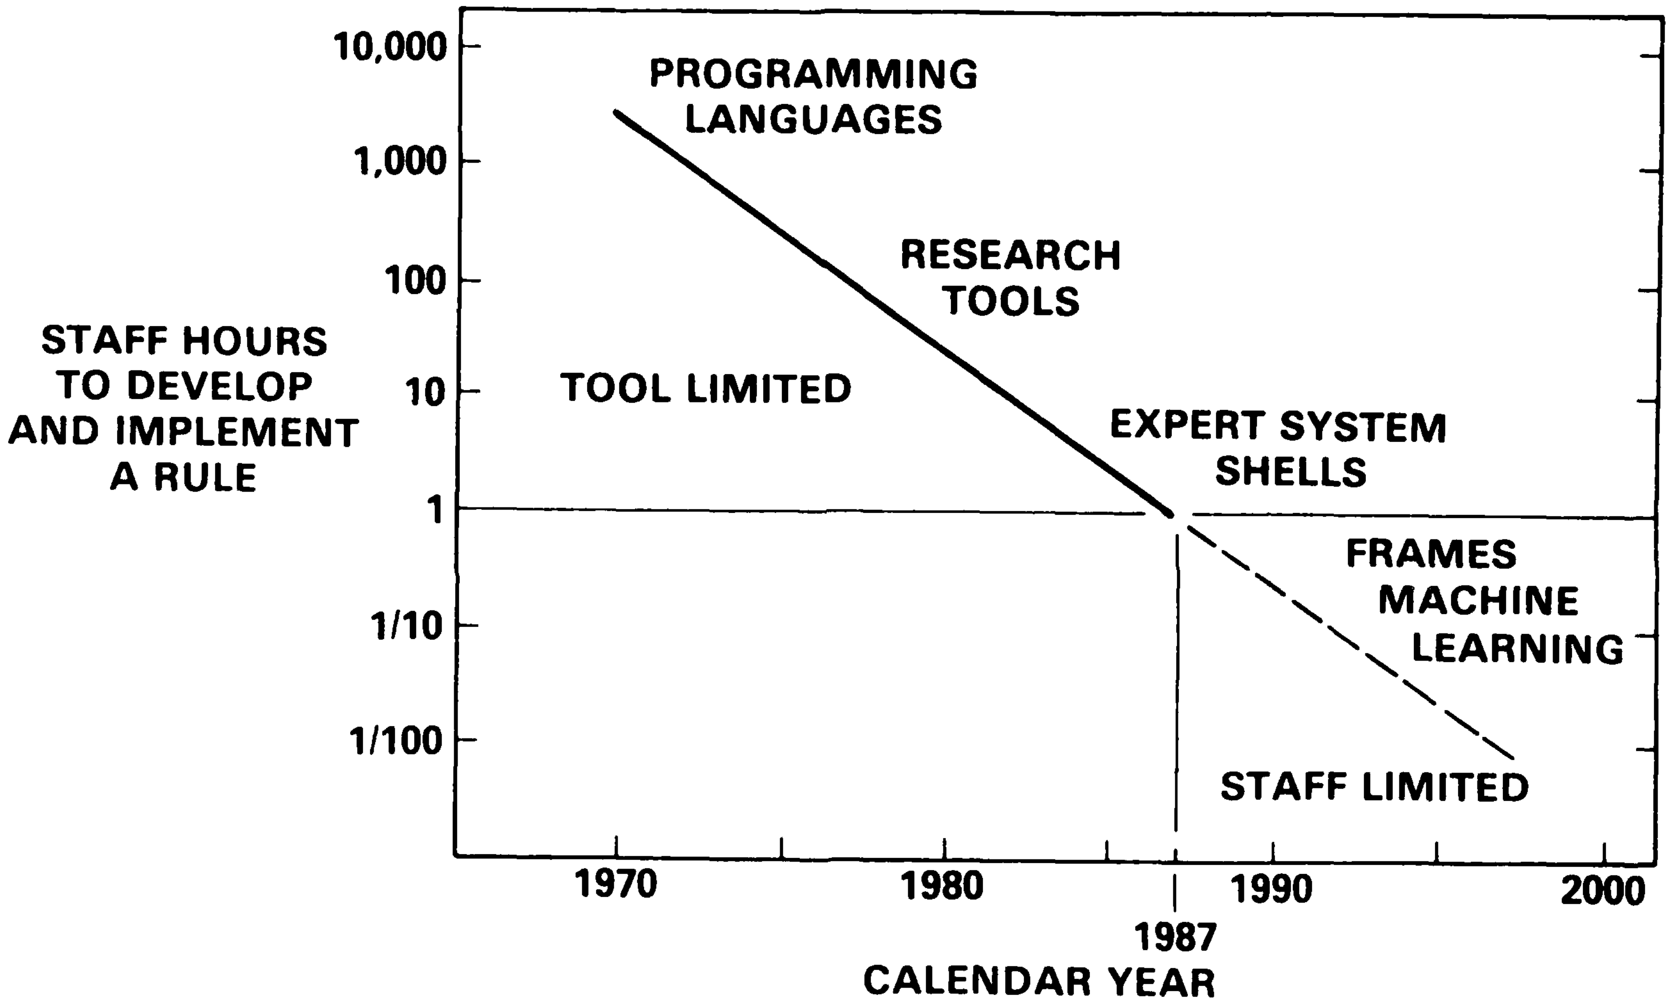

Figure 1 - Expert System Development Trends and Projections

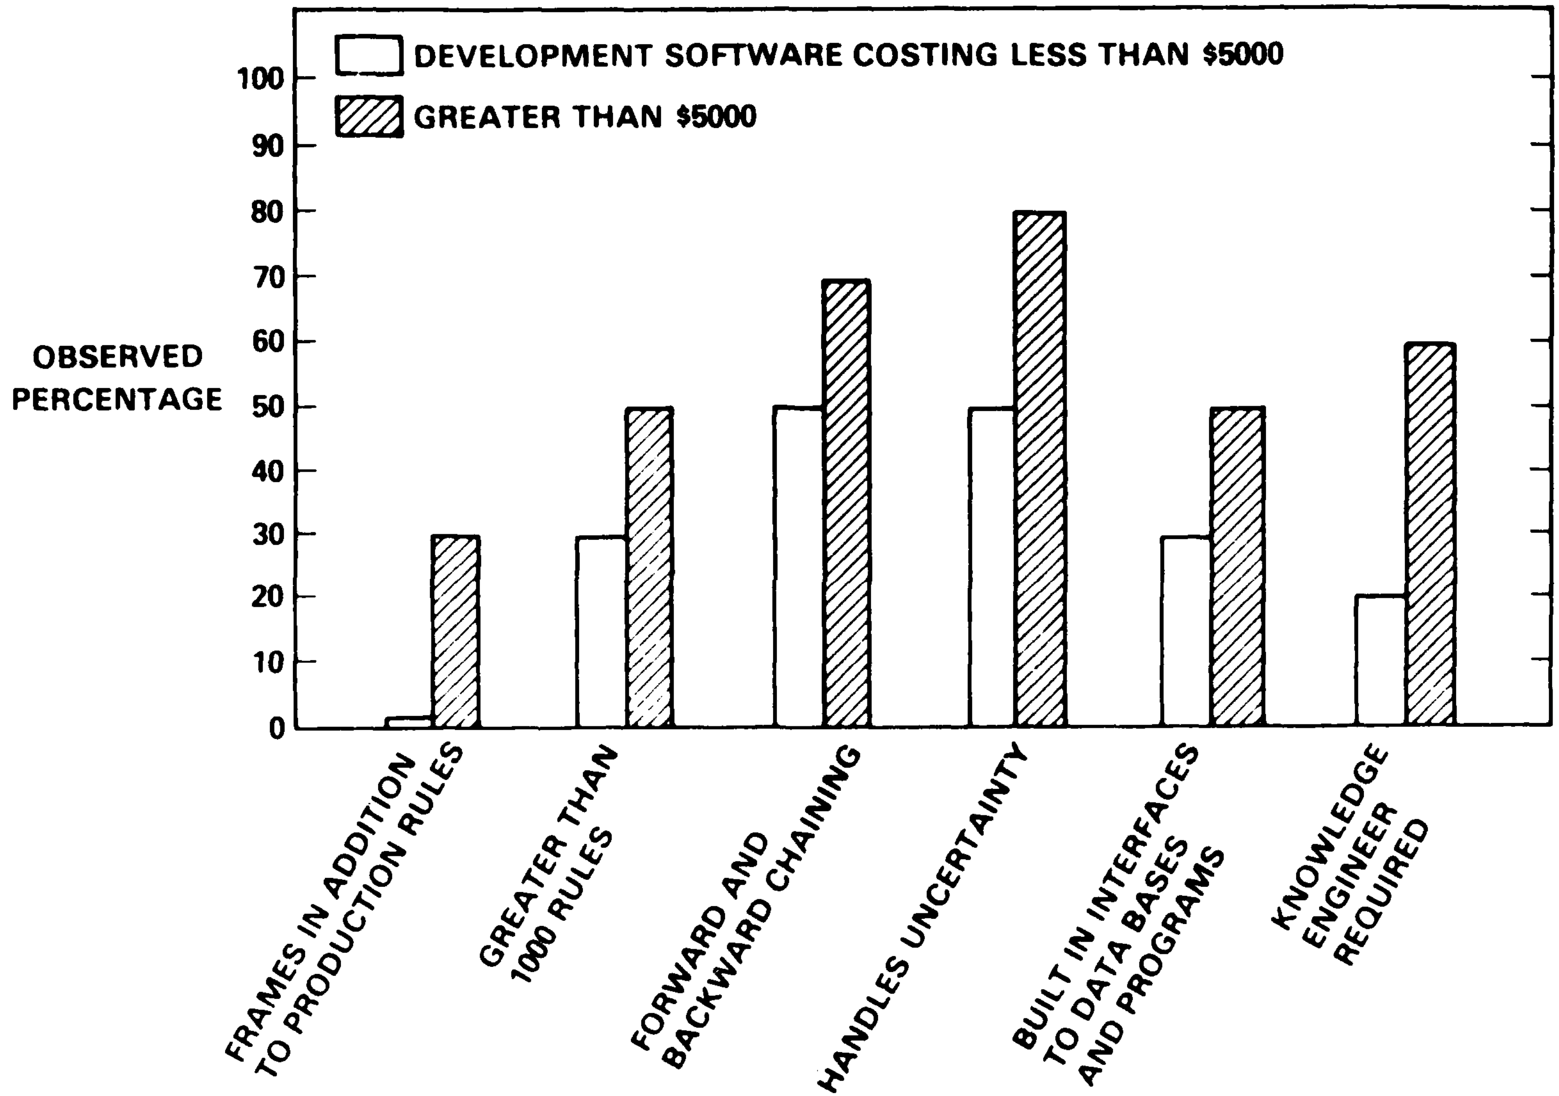

Figure 2 - Key Differences Between Expert Systems

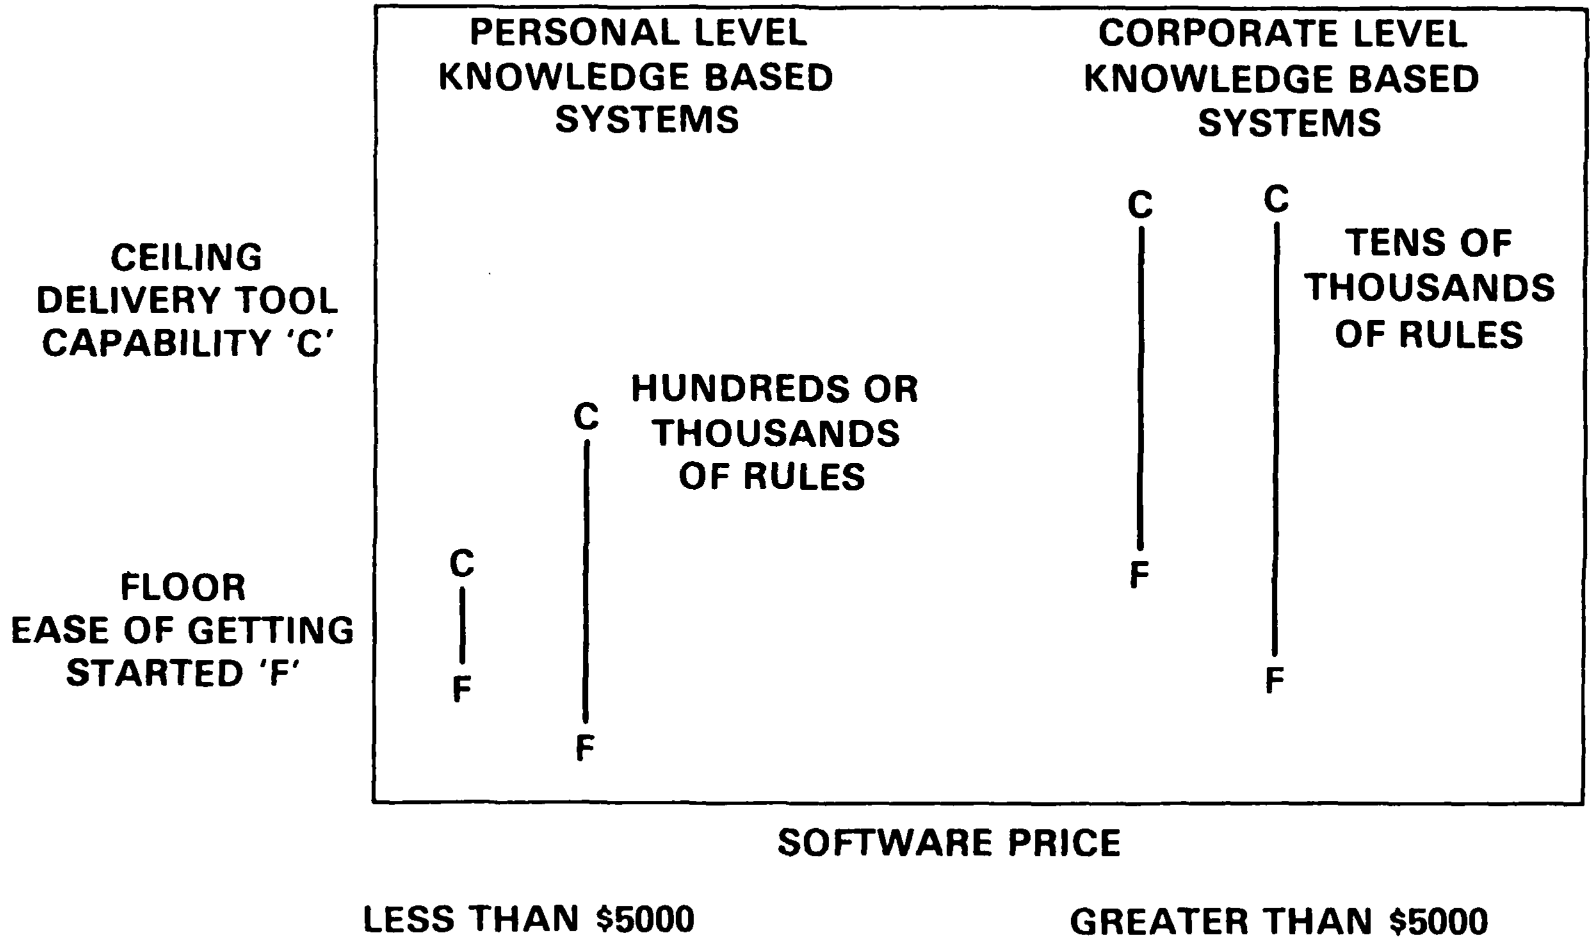

Figure 3 - The Concept of 'Floors' and 'Ceilings' in Expert Systems

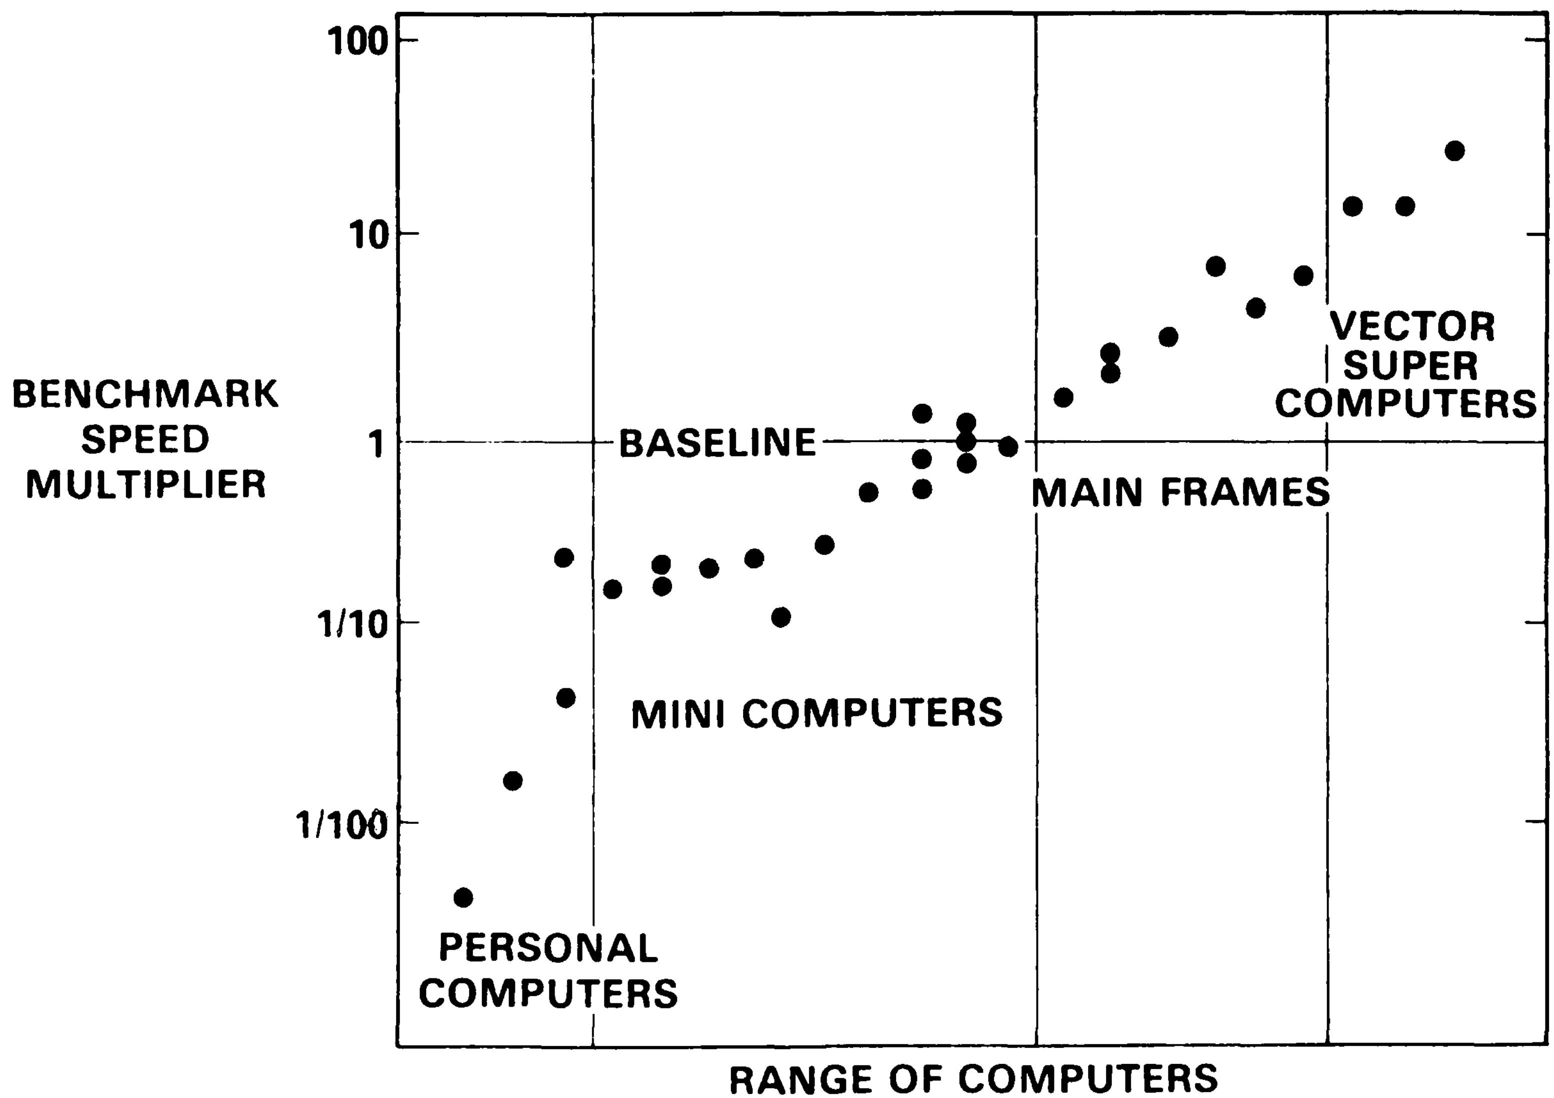

Figure 4 - Relative Machine Performance

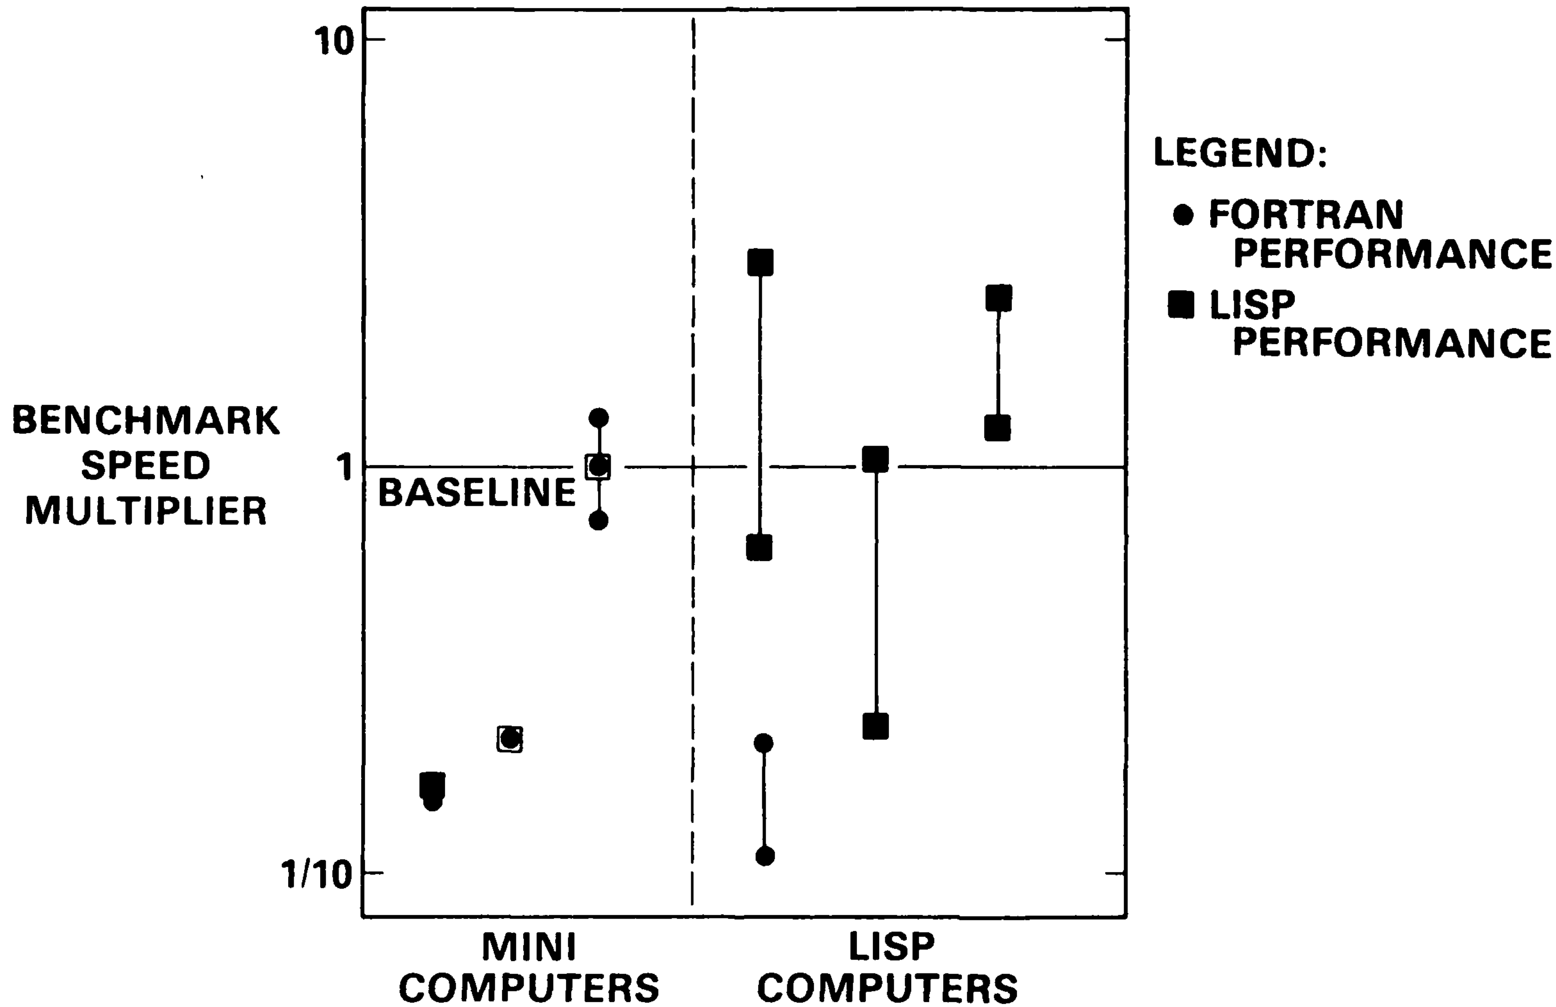

Figure 5 - FORTRAN and LISP Performance

CORPORATE LEVEL
$300,000 — 1,500,000

PERSONAL LEVEL
$35,000 — 70,000

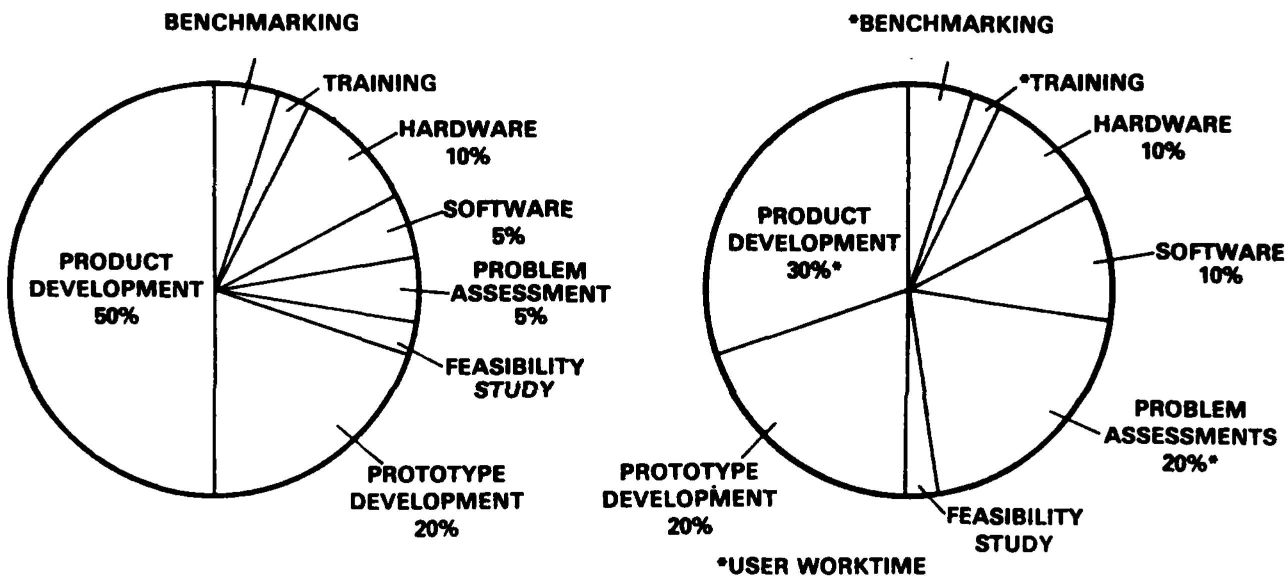

YEARLY MAINTENANCE COSTS:

80% OF INITIAL SYSTEM COSTS
(3-12 PEOPLE PER YEAR)

YEARLY MAINTENANCE COSTS

30%* OF INITIAL SYSTEM COSTS

Figure 6 - Distribution of Expert System Development Costs

Table 1 - Some of the Areas Appropriate to Artificial Intelligence

ADAPTIVE SYSTEMS	KNOWLEDGE ACQUISITION	QUALITY CONTROL
ARTIFICIAL INTELLIGENCE	KNOWLEDGE BASED SYSTEMS	REASONING/DECISION MAKING
AUTOMATIC PROGRAMMING	KNOWLEDGE REPRESENTATION	MAKING RECONNAISSANCE
AUTOMATIC REPAIR	LOGISTICS	RESOURCE MANAGEMENT
AUTONOMOUS VEHICLES	MACHINE LEARNING	ROBOTICS
CAD	MAINTENANCE	SENSORS
COMMAND AND CONTROL	MANAGEMENT	SHIPBUILDING
COMMUNICATIONS	MAN MACHINE SYSTEMS	SIGNAL PROCESSING
COMPUTER BASED INSTRUCTION	MANUFACTURING TECHNOLOGY	SIMULATION
COMPUTER HARDWARE	MATERIALS HANDLING	SPEECH RECOGNITION
COMPUTER PROGRAMS	NATURAL LANGUAGE PROCESSING	STRATEGY ANALYSIS
CONFERENCES	NAVIGATION	TACTICAL PROBLEM SOLVING
COUNTERMINE VEHICLE	NIGHT VISION	TARGET ANALYSIS
DATA FUSION	OPERATIONS RESEARCH	TERRAIN ANALYSIS
DATABASE MANAGEMENT	PAINTING	THREAT EVALUATION
DECISION AIDS	PATTERN RECOGNITION	TRAINING
ELECTRONIC WARFARE	PERFORMANCE MEASUREMENT	VISION SYSTEMS
ENERGY CONSERVATION	PLANNING	WEAPON SYSTEMS
EXPERT SYSTEMS	POWER SOURCES	WELDING
IMAGE PROCESSING	PROBLEM SOLVING	
INFORMATION FUSION	PROGRAMMING TOOLS	
INSTRUMENTATION		

Table 2 - Where Expert Systems Fit

GAGE	NUMBER OF RULES
• THE LIMIT OF HUMAN EXPERTISE	100,000
• EXPERTISE IN A PROFESSION	10,000
• EXPERT COMPETENCE IN A NARROW AREA	500 — 1000
• CONVINCING EXPERT SYSTEM DEMONSTRATION	250
• COMMERCIALLY PRACTICAL EXPERT SYSTEM	50

Table 3 - Current Versus Future Areas of AI Application

• CLASSIFICATION	STRUCTURED SELECTION PROCESS. DIAGNOSIS IN A SPECIALITY. SITUATION ANALYSIS.	HIGHLY SUCCESSFUL (NOW)
• DESIGN SYNTHESIS	DESIGN BY PLAN SELECTION AND REFINEMENT	SUCCESSFUL (NOW)
• KNOWLEDGE BASED DATA STORAGE AND RETREIVAL	CORRELATE ABSTRACT DATA. COMPLEX KNOWLEDGE BASE.	NEAR TERM (LIKELY SOON)
• ABDUCTIVE ASSEMBLY OF HYPOTHESES	GENERATE AND ASSEMBLY OF HYPOTHESES.	MID TERM (POSSIBLE)
• ABSTRACTION	"WHAT WILL HAPPEN IF. . .?" REASONING OF EXPERTS. GENERAL, VARIED, AND WIDE RANGING.	FAR TERM (NEXT GENERATION)
• HYPOTHESIS MATCHING	MATCHING HYPOTHESES AND DATA.	FAR TERM (NEXT GENERATION)
• BROAD EXPERTISE	COMMON SENSE QUALITATIVE REASONING INVOLVING TIME AND SPACE	FAR TERM (NEXT GENERATION)

Table 4 - Absolute Needs for Expert Sytem Development

• RULES AND FRAMES	– 1000 RULES +
• SPEED	– WRITTEN IN LISP/C FOR DEVELOPMENT SOFTWARE
	– C FOR RUNTIME SOFTWARE
• TRANSPORTABLE	– YES, MUST RUN ON NUMEROUS DEVELOPMENT AND DELIVERY MACHINES WITHOUT MAJOR CHANGE
• EASY TO USE	– YES, NATURAL LANGUAGE INTERFACE AND MENUS
• HANDLE UNCERTAINTY	– YES, FUZZY LOGIC AND INCOMPLETE INFORMATION
• POWERFUL INTERFACES	– YES, TO PROGRAMS, DATA BASES, AND FUNCTIONS
• EXPLAINATIONS	– YES
• KNOWLEDGE ENGINEER NEEDED	– DEPENDS ON PROJECT COMPLEXITY
• PLAIN ENGLISH PROGRAMMING	– DEPENDS ON PROJECT COMPLEXITY
• PRICE	– DEPENDS ON PROJECT COMPLEXITY
• MACHINE	– DEPENDS ON PROJECT COMPLEXITY

Table 5 - Distribution of Expert System Development Costs

AREA	CORPORATE LEVEL KNOWLEDGE BASED SYSTEMS	PERSONAL LEVEL KNOWLEDGE BASED SYSTEMS
• HARDWARE	10% ($50,000 – 100,000)	10% ($5,000 – 10,000)
• SOFTWARE	5% ($5,000 – 90,000)	10% ($3,000 – 5,000)
• PROBLEM ASSESSMENT	5%	20%*
• FEASIBILITY STUDY	<5%	<5%
• PROTOTYPE SYSTEM	20%	20%
• PRODUCT DEVELOPMENT	50%	30%*
• BENCHMARKING	5%	5%*
• TRAINING	<5%	<5%*
TOTAL	$300,000 – 1,500,000	$35,000 – 70,000
• YEARLY MAINTENANCE	80% OF INITIAL SYSTEM COSTS (3 – 12 PEOPLE)	30%* OF INITIAL SYSTEM COSTS

*USER WORKTIME

END
9-87
DTIC